

ATTORNEY'S DOCKET NO. 06502.0381  
SUN REFERENCE NO. P5189  
PATENTS

VIA EXPRESS MAIL - EL396216485US

**UNITED STATES PATENT APPLICATION**

**OF**

**GUY L. STEELE JR.**

**FOR**

**FLOATING POINT MULTIPLIER WITH EMBEDDED STATUS INFORMATION**

06502.0381

## RELATED APPLICATIONS

[001] U.S. Patent Application Serial No. \_\_\_\_\_, filed on even date herewith in the name of Guy L. Steele Jr. and entitled "Floating Point System That Represents Status Flag Information Within a Floating Point Operand," assigned to the assignee of the present application, is hereby incorporated by reference.

## DESCRIPTION OF THE INVENTION

### Field of the Invention

[002] The invention relates generally to systems and methods for performing floating point operations, and more particularly to systems and methods for performing floating point multiplication with embedded status information associated with a floating point operand.

### Background of the Invention

[003] Digital electronic devices, such as digital computers, calculators and other devices, perform arithmetic calculations on values in integer, or "fixed point," format, in fractional, or "floating point" format, or both. Institute of Electrical and Electronic Engineers (IEEE) Standard 754, (hereinafter "IEEE Std. 754" or "the Standard") published in 1985 and adopted by the American National Standards Institute (ANSI), defines several standard formats for expressing values in floating point format and a number of aspects regarding behavior of computation in connection therewith. In accordance with IEEE Std. 754, a representation in floating point format comprises a plurality of binary digits, or "bits," having the structure

[004]  $s e_{msb} \cdots e_{lsb} f_{msb} \cdots f_{lsb}$

where bit “s” is a sign bit indicating whether the entire value is positive or negative, bits “ $e_{msb} \cdots e_{lsb}$ ” comprise an exponent field that represents the exponent “e” in unsigned binary biased format, and bits “ $f_{msb} \cdots f_{lsb}$ ” comprise a fraction field that represents the fractional portion “f” in unsigned binary format (“msb” represents “most significant bit” and “lsb” represents “least significant bit”). The Standard defines two general formats. A “single” format comprises thirty-two bits while a “double” format comprises sixty-four bits. In the single format, there is one sign bit “s,” eight bits “ $e_7 \dots e_0$ ” comprising the exponent field and twenty-three bits “ $f_{22} \dots f_0$ ” comprising the fraction field. In the double format, there is one sign bit “s,” eleven bits “ $e_{10} \dots e_0$ ” comprising the exponent field and fifty-two bits “ $f_{51} \dots f_0$ ” comprising the fraction field.

[005] As indicated above, the exponent field of the floating point representation “ $e_{msb} \cdots e_{lsb}$ ” represents the exponent “E” in biased format. The biased format provides a mechanism by which the sign of the exponent is implicitly indicated. In particular, the bits “ $e_{msb} \cdots e_{lsb}$ ” represent a binary encoded value “e” such that “ $e = E + \text{bias}$ .” This allows the exponent E to extend from -126 to +127, in the eight-bit “single” format, and from -1022 to +1023 in the eleven-bit “double” format, and provides for relatively easy manipulation of the exponents in multiplication and division operations, in which the exponents are added and subtracted, respectively.

[006] IEEE Std. 754 provides for several different formats with both the single and double formats which are generally based on the bit patterns of the bits

" $e_{msb} \cdots e_{lsb}$ " comprising the exponent field and the bits " $f_{msb} \cdots f_{lsb}$ " comprising the fraction field. If a number is represented such that all of the bits " $e_{msb} \cdots e_{lsb}$ " of the exponent field are binary one's (i.e., if the bits represent a binary-encoded value of "255" in the single format or "2047" in the double format) and all of the bits " $f_{msb} \cdots f_{lsb}$ " of the fraction field are binary zeros, then the value of the number is positive or negative infinity, depending on the value of the sign bit "s." In particular, the value "v" is  $v = (-1)^s \infty$ , where " $\infty$ " represents the value "infinity." On the other hand, if all of the bits " $e_{msb} \cdots e_{lsb}$ " of the exponent field are binary one's and if the bits " $f_{msb} \cdots f_{lsb}$ " of the fraction field are not all zero's, then the value that is represented is deemed "not a number," which is abbreviated in the Standard by "NaN."

[007] If a number has an exponent field in which the bits " $e_{msb} \cdots e_{lsb}$ " are neither all binary ones nor all binary zeros (i.e., if the bits represent a binary-encoded value between 1 and 254 in the single format or between 1 and 2046 in the double format), the number is said to be a "normalized" format. For a number in the normalized format, the value represented by the number is

$$v = (-1)^s 2^{e-bias} (1.f_{msb} \cdots f_{lsb}), \text{ where "||" represents a concatenation operation.}$$

Effectively, in the normalized format, there is an implicit most significant digit having the value "one," so that the twenty-three digits in the fraction field of the single format, or the fifty-two digits in the fraction field of the double format, will effectively represent a value having twenty-four digits or fifty-three digits of precision, respectively, where the value is less than two, but not less than one.

[008] On the other hand, if a number has an exponent field in which the bits " $e_{msb} \cdots e_{lsb}$ " are all binary zeros, representing the binary-encoded value of "zero," and

a fraction field in which the bits  $f_{msb} \cdots f_{lsb}$  are not all zero, the number is said to be a “de-normalized” format. For a number in the de-normalized format, the value represented by the number is  $v = (-1)^s 2^{e-bias+1} (0.f_{msb} \cdots f_{lsb})$ . It will be appreciated that the range of values of numbers that can be expressed in the de-normalized format is disjoint from the range of values of numbers that can be expressed in the normalized format, for both the single and double formats. Finally, if a number has an exponent field in which the bits “ $e_{msb} \cdots e_{lsb}$ ” are all binary zeros, representing the binary-encoded value of “zero,” and a fraction field in which the bits  $f_{msb} \cdots f_{lsb}$  are all zero, the number has the value “zero”. It will be appreciated that the value “zero” may be positive zero or negative zero, depending on the value of the sign bit.

[009] Generally, circuits or devices that perform floating point computations or operations (generally referred to as floating point units) conforming to IEEE Std. 754 are designed to generate a result in three steps:

[010] (a) In the first step, an approximation calculation step, an approximation to the absolutely accurate mathematical result (assuming that the input operands represent the specific mathematical values as described by IEEE Std. 754) is calculated that is sufficiently precise as to allow this accurate mathematical result to be summarized. The summarized result is usually represented by a sign bit, an exponent (typically represented using more bits than are used for an exponent in the standard floating-point format), and some number “N” of bits of the presumed result fraction, plus a guard bit and a sticky bit. The value of the exponent will be such that the value of the fraction generated in step (a) consists of a 1 before the binary point and a fraction after the binary point. The bits

are commonly calculated so as to obtain the same result as the following conceptual procedure (which is impossible under some circumstances to carry out in practice): calculate the mathematical result to an infinite number of bits of precision in binary scientific notation, and in such a way that there is no bit position in the significand such that all bits of lesser significance are 1-bits (this restriction avoids the ambiguity between, for example, 1.100000... and 1.011111... as representations of the value "one-and-one-half"); let the N most significant bits of the infinite significand be used as the intermediate result significand; let the next bit of the infinite significand be the guard bit; and let the sticky bit be 0 if and only if ALL remaining bits of the infinite significand are 0-bits (in other words, the sticky bit is the logical OR of all remaining bits of the infinite fraction after the guard bit).

[011] (b) In the second step, a rounding step, the guard bit, the sticky bit, perhaps the sign bit, and perhaps some of the bits of the presumed significand generated in step (a) are used to decide whether to alter the result of step (a). For conventional rounding modes defined by IEEE Std. 754, this is a decision as to whether to increase the magnitude of the number represented by the presumed exponent and fraction generated in step (a). Increasing the magnitude of the number is done by adding 1 to the significand in its least significant bit position, as if the significand were a binary integer. It will be appreciated that, if the significand is all 1-bits, then the magnitude of the number is "increased" by changing it to a high-order 1-bit followed by all 0-bits and adding 1 to the exponent.

[012] Regarding the rounding modes, it will be further appreciated that,

[013] (i) if the result is a positive number, and

[014] (a) if the decision is made to increase, effectively the decision has been made to increase the value of the result, thereby rounding the result up (*i.e.*, towards positive infinity), but

[015] (b) if the decision is made not to increase, effectively the decision has been made to decrease the value of the result, thereby rounding the result down (*i.e.*, towards negative infinity); and

[016] (ii) if the result is a negative number, and

[017] (a) if the decision is made to increase, effectively the decision has been made to decrease the value of the result, thereby rounding the result down, but

[018] (b) if the decision is made not to increase, effectively the decision has been made to increase the value of the result, thereby rounding the result up.

[019] (c) In the third step, a packaging step, a result is packaged into a standard floating-point format. This may involve substituting a special representation, such as the representation defined for infinity or NaN if an exceptional situation (such as overflow, underflow, or an invalid operation) was detected. Alternatively, this may involve removing the leading 1-bit (if any) of the fraction, because such leading 1-bits are implicit in the standard format. As another alternative, this may involve shifting the fraction in order to construct a denormalized number. As a specific example, it is assumed that this is the step that forces the result to be a NaN if any input operand is a NaN. In this step, the decision is also made as to whether the result should be an infinity. It will be appreciated that, if the

result is to be a NaN or infinity from step (b), the original result will be discarded and an appropriate representation will be provided as the result.

[020] In addition in the packaging step, floating point status information is generated, which is stored in a floating point status register. The floating point status information generated for a particular floating point operation includes indications, for example, as to whether

[021] (i) a particular operand is invalid for the operation to be performed (“invalid operation”);

[022] (ii) if the operation to be performed is division, the divisor is zero (“division-by-zero”);

[023] (iii) an overflow occurred during the operation (“overflow”);

[024] (iv) an underflow occurred during the operation (“underflow”);

and

[025] (v) the rounded result of the operation is not exact (“inexact”).

[026] These conditions are typically represented by flags that are stored in the floating point status register. The floating point status information can be used to dynamically control the operations in response to certain instructions, such as conditional branch, conditional move, and conditional trap instructions that may be in the instruction stream subsequent to the floating point instruction. Also, the floating point status information may enable processing of a trap sequence, which will interrupt the normal flow of program execution. In addition, the floating point status information may be used to affect certain ones of the functional unit control



signals that control the rounding mode. IEEE Std. 754 also provides for accumulating floating point status information from, for example, results generated for a series or plurality of floating point operations.

[027] IEEE Std. 754 has brought relative harmony and stability to floating-point computation and architectural design of floating-point units. Moreover, its design was based on some important principles, and rests on a sensible mathematical semantics that eases the job of programmers and numerical analysts. It also supports the implementation of interval arithmetic, which may prove to be preferable to simple scalar arithmetic for many tasks. Nevertheless, IEEE Std. 754 has some serious drawbacks, including:

[028] (i) Modes (e.g., the rounding modes and traps enabled/disabled mode), flags (e.g., flags representing the status information), and traps required to implement IEEE Std. 754 introduce implicit serialization issues. Implicit serialization is essentially the need for serial control of access (read/write) to and from globally used registers, such as a floating point status register. Under IEEE Std. 754, implicit serialization may arise between (1) different concurrent floating-point instructions and (2) between floating point instructions and the instructions that read and write the flags and modes. Furthermore, rounding modes may introduce implicit serialization because they are typically indicated as global state, although in some microprocessor architectures, the rounding mode is encoded as part of the instruction operation code, which will alleviate this problem to that extent. Thus, the potential for implicit serialization makes the Standard difficult

to implement coherently and efficiently in today's superscalar and parallel processing architectures without loss of performance.

[029] (ii) The implicit side effects of a procedure that can change the flags or modes can make it very difficult for compilers to perform optimizations on floating point code. As a result, compilers for most languages usually assume that every procedure call is an optimization barrier in order to be safe. This unfortunately may lead to further loss of performance.

[030] (iii) Global flags, such as those that signal certain modes, make it more difficult to do instruction scheduling where the best performance is provided by interleaving instructions of unrelated computations. Thus, instructions from regions of code governed by different flag settings or different flag detection requirements cannot easily be interleaved when they must share a single set of global flag bits.

[031] (iv) Furthermore, traps have been difficult to integrate efficiently into computing architectures and programming language designs for fine-grained control of algorithmic behavior.

[032] Thus, there is a need for a system that avoids such problems when performing floating point operations and, in particular, when performing floating point multiplication with embedded status information associated with a floating point operand.

## SUMMARY OF THE INVENTION

[033] Consistent with the current invention, a floating point multiplier with embedded status information method and system are provided that avoid the

problems associated with prior art floating point multiplier systems as discussed herein above.

[034] In one aspect, a system for providing a floating point product comprises an analyzer circuit configured to determine a first status of a first floating point operand and a second status of a second floating point operand based upon data within the first floating point operand and the second floating point operand respectively. In addition, the system comprises a results circuit coupled to the analyzer circuit. The results circuit is configured to assert a resulting floating point operand containing the product of the first floating point operand and the second floating point operand and a resulting status embedded within the resulting floating point operand.

[035] In another aspect, a method for providing a floating point product comprises determining a first status of a first floating point operand and a second status of a second floating point operand based upon data within the first floating point operand and the second floating point operand respectively. In addition, the method comprises asserting a resulting floating point operand containing the product of the first floating point operand and the second floating point operand and a resulting status embedded within the resulting floating point operand.

[036] In yet another aspect, a computer-readable medium on which is stored a set of instructions for providing a floating point product, which when executed perform stages comprising determining a first status of a first floating point operand and a second status of a second floating point operand based upon data within the first floating point operand and the second floating point operand respectively. The

instruction set further comprises asserting a resulting floating point operand containing the product of the first floating point operand and the second floating point operand and a resulting status embedded within the resulting floating point operand.

[037] Both the foregoing general description and the following detailed description are exemplary and are intended to provide further explanation of the invention as claimed.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[038] The accompanying drawings provide a further understanding of the invention and, together with the detailed description, explain the principles of the invention. In the drawings:

[039] FIG. 1 is a functional block diagram of an exemplary system for providing a floating point product consistent with an embodiment of the present invention;

[040] FIG. 2 illustrates exemplary formats for representations of floating point values generated by the system of FIG. 1 consistent with an embodiment of the present invention;

[041] FIG. 3 illustrates a table useful in understanding the operations of the exemplary system of FIG. 1 consistent with an embodiment of the present invention; and

## **DESCRIPTION OF THE EMBODIMENTS**

[illegible][illegible][illegible]

result and, in some cases, floating point status information, with the floating point status information being encoded in and comprising part of the floating point representation of the result. Since the floating point status information comprises part of the floating point representation of the result, instead of being separate and apart from the result as in prior art multiplier units, the implicit serialization that is required by maintaining the floating point status information separate and apart from the result can be obviated.

[046] The multiplier unit 10 encodes the floating point status information in results that are generated in certain formats. This will be illustrated in connection with FIG. 2. FIG. 2 depicts exemplary formats of floating point operands that the multiplier unit 10 may receive and of results that it generates. With reference to the embodiment illustrated in FIG. 2, seven formats are depicted, including a zero format 70, an underflow format 71, a denormalized format 72, a normalized non-zero format 73, an overflow format 74, an infinity format 75 and a not-a-number (NaN) format 76. The zero format 70 is used to represent the values "zero," or, more specifically, positive or negative zero, depending on the value of "s," the sign bit.

[047] The underflow format 71 provides a mechanism by which the multiplier unit 10 can indicate that the result of a computation is an underflow. In the underflow format, the sign bit "s" indicates whether the result is positive or negative, the bits  $e_{msb} \dots e_{lsb}$  of the exponent field are all binary zero's, and the bits  $f_{msb} \dots f_{lsb+1}$  of the fraction field, except for the least significant bit, are all binary zero's. The least significant bit  $f_{lsb}$  of the fraction field is a binary one.

[048] The denormalized format 72 and normalized non-zero format 73 are used to represent finite non-zero floating point values substantially along the lines of that described above in connection with IEEE Std. 754. In both formats 72 and 73, the sign bit "s" indicates whether the result is positive or negative. The bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the denormalized format 72 are all binary zero's, whereas the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the normalized non-zero format 73 are mixed one's and zero's, except that the exponent field of the normalized non-zero format 73 will not have a pattern in which bits  $e_{msb} \dots e_{lsb+1}$  are all binary ones and the least significant bit  $e_{lsb}$  zero and all of the bits  $f_{msb} \dots f_{lsb}$  of the fraction field are all binary one's. In both formats 72 and 73, the bits  $f_{msb} \dots f_{lsb}$  of the fraction field are not all binary zero's.

[049] The overflow format 74 provides a mechanism by which the multiplier unit 10 can indicate that the result of a computation is an overflow. In the overflow format 74, the sign bit "s" indicates whether the result is positive or negative, the bits  $e_{msb} \dots e_{lsb+1}$  of the exponent field are all binary ones, with the least significant bit  $e_{lsb}$  being zero. The bits  $f_{msb} \dots f_{lsb}$  of the fraction field are all binary ones.

[050] The infinity format 75 provides a mechanism by which the multiplier unit 10 can indicate that the result is infinite. In the infinity format 75, the sign bit "s" indicates whether the result is positive or negative, the bits  $e_{msb} \dots e_{lsb}$  of the exponent field are all binary ones, and the bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field are all binary zero's. The five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field are flags, which will be described below.

[051] The NaN format 76 provides a mechanism by which the multiplier unit 10 can indicate that the result is not a number. In the NaN format, the sign bit "s" can be any value, the bits  $e_{msb} \dots e_{lsb}$  of the exponent field are all binary ones, and the bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field are not all binary zero's. The five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field are flags, which will be described below.

[052] As noted above, in values represented in the infinity format 75 and the NaN format 76, the five low order bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field are flags. In the formats used with the multiplier unit 10, the five flags include the flags that are defined by IEEE Std. 754, including an invalid operation flag "n", an overflow flag "o", an underflow flag "u", a division-by-zero flag "z", and an inexact flag "x". For example, a value in the NaN format 76 in which both the overflow flag "o" and the division-by-zero flag "z" are set, indicates that the value represents a result of a computation that involved an overflow (this from the overflow flag "o"), as well as an attempt to divide by zero (this from the division-by-zero flag "z"). It should be noted that the flags provide the same status information as provided by, for example, information stored in a floating point status register in a prior art floating point unit. Because the information is provided as part of the result and stored therewith in registers in which the result is ultimately stored, multiple instructions can be contemporaneously executed. This is because the floating point status information that may be generated during execution of one instruction, when stored, will not over-write previously-stored floating point status information generated during execution of another instruction.



[053] In addition to including status information in the five low-order bits  $f_{lsb+4}...f_{lsb}$  of the fraction field for values in the NaN format 76, other information can also be encoded in the next five low-order bits  $f_{lsb+9}...f_{lsb+5}$ . If the value in the NaN format 76 is the result of an operation, the other information indicates the operation and types of operands that gave rise to the result. In one embodiment, the other information is associated with binary encoded values (BEV) of those bits  $f_{lsb+9}...f_{lsb+5}$  as follows:

Bit Pattern Of Result	BEV of $f_{lsb+9}...f_{lsb+5}$	Meaning
	0 or 1	no specific meaning
s 11111111 00000000000000010nouzx	2	infinity minus infinity
s 11111111 00000000000000011nouzx	3	OV minus OV
s 11111111 000000000000000100nouzx	4	zero times infinity
s 11111111 000000000000000101nouzx	5	UN times OV
	6 or 7	no specific meaning
s 11111111 00000000000001000nouzx	8	zero divided by zero
s 11111111 00000000000001001nouzx	9	infinity divided by infinity
s 11111111 00000000000001010nouzx	10	UN divided by UN
s 11111111 00000000000001011nouzx	11	OV divided by OV
s 11111111 00000000000001100nouzx	12	square root of less than zero
	13-16	no specific meaning
s 11111111 000000000000010001nouzx	17	remainder by zero
s 11111111 000000000000010010nouzx	18	remainder by UN
s 11111111 000000000000010011nouzx	19	remainder by OV
s 11111111 000000000000010100nouzx	20	remainder of infinity
s 11111111 000000000000010101nouzx	21	remainder of infinity by zero
s 11111111 000000000000010110nouzx	22	remainder of infinity by UN
s 11111111 000000000000010111nouzx	23	remainder of infinity by OV
s 11111111 000000000000011000nouzx	24	remainder of OV
s 11111111 000000000000011001nouzx	25	remainder of OV by zero
s 11111111 000000000000011010nouzx	26	remainder of OV by UN
s 11111111 000000000000011011nouzx	27	remainder of OV by OV
	28-31	no specific meaning

[054] In the following, it will be assumed that the formats represent thirty-two bit values; extension to, for example, sixty-four bit values or values represented in other numbers of bit will be readily apparent to those skilled in the art. Additionally,

"OV" refers to an operand in the overflow format 74, "UN" refers to an operand in the underflow format 71 and "infinity" refers to an operand in the infinity format 75.

[055] In addition, it will be convenient in the following to have names for the finite nonzero numbers that are adjacent to +OV (a value in the overflow pattern with the sign bit "s" having the value "zero" indicating a positive value), -OV (a value in the overflow pattern with the sign bit "s" having the value "one," indicating a negative value), +UN (a value in the underflow pattern with the sign bit "s" having the value "zero," indicating a positive value), and -UN (a value in the underflow pattern with the sign bit "s" having the value "one," indicating a negative value), as follows:

0 00000000 000000000000000000000010	+TINY
1 00000000 000000000000000000000010	-TINY
0 11111110 111111111111111111111110	+HUGE
1 11111110 111111111111111111111110	-HUGE

[056] Generally, +OV can be deemed to refer to "some (or any) value that is strictly between +HUGE and  $+\infty$ " and +UN can be deemed to refer to "some (or any) value that is strictly between +0 and +TINY". Similarly, -OV can be deemed to refer to "some (or any) value that is strictly between -HUGE and  $-\infty$ " and -UN can be deemed to refer to "some (or any) value that is strictly between -0 and -TINY."

[057] With this background, the structure and operation of the exemplary multiplier unit 10 will be described in connection with FIG. 1 and consistent with an embodiment of the invention. With reference to FIG. 1, the exemplary multiplier unit 10 includes two operand buffers 11A and 11B, respective operand analysis circuits 12A and 12B, a multiplier core 13, a result assembler 14 and a multiplier decision table logic circuit 15. The operand buffers 11A and 11B receive and store

respective operands from, for example, a set of registers (not shown) in a conventional manner. The multiplier core 13 receives the operands from the operand buffers 11A and 11B, except as described below, and rounding mode information from, for example, a rounding mode store 16. The multiplier core 13 then generates a result in accordance with IEEE Std. 754. Multiplier core 13 is conventional and will not be described in detail herein.

[058] Each operand analysis circuit 12A, 12B analyzes the operand in the respective buffer 11A, 11B and generates signals providing information relating thereto, which signals are provided to the multiplier decision table logic circuit 15. The result assembler 14 receives information from a number of sources, including the operand buffers 11A and 11B, multiplier core 13 and several predetermined value stores as described below. Under control of control signals from the multiplier decision table logic circuit 15, the result assembler 14 assembles the result, which is provided on a result bus 17. The result bus 17, in turn, may deliver the result to any convenient destination, such as a register in a register set (not shown), for storage or other use.

[059] The system for providing a floating point product may comprise an analyzer circuit configured to determine a first status of a first floating point operand and a second status of a second floating point operand based upon data within the first floating point operand and the second floating point operand, respectively. In one embodiment, the analyzer circuit includes buffers 11A, 11B and analysis circuits 12A, 12B. In addition, the system for providing a floating point product includes a results circuit coupled to the analyzer circuit. The results circuit is configured to

assert a resulting floating point operand containing the product of the first floating point operand and the second floating point operand and a resulting status embedded within the resulting floating point operand. The results circuit may be implemented with a multiplier circuit (comprising the multiplier core 13), the multiplier decision logic table circuit 15, and result assembler 14.

[060] Those skilled in the art will appreciate that the invention may be practiced in an electrical circuit comprising discrete electronic elements, packaged or integrated electronic chips containing logic gates, a circuit utilizing a microprocessor, or on a single chip containing electronic elements or microprocessors. It may also be provided using other technologies capable of performing logical operations such as, for example, AND, OR, and NOT, including but not limited to mechanical, optical, fluidic, and quantum technologies. In addition, the invention may be practiced within a general purpose computer or in any other circuits or systems as are known by those skilled in the art.

[061] As noted above, each operand analysis circuit 12A, 12B analyzes the operand in the respective buffer 11A, 11B and generates signals providing information relating thereto. These signals are provided to the multiplier decision table logic circuit 15. In the exemplary embodiment, each operand analysis circuit 12A, 12B is implemented with a number of comparators, including:

[062] (i) a comparator 20A, 20B that generates an asserted signal if the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in respective buffer 11A, 11B are all binary one's, which will be the case if the operand is in the infinity format 75 or the NaN format 76;

[063] (ii) a comparator 21A, 21B that generates an asserted signal if the bits  $e_{msb} \dots e_{lsb+1}$  of the exponent field of the operand in the respective buffer 11A, 11B are all binary one's and the bit  $e_{lsb}$  is a binary zero, which will be the case if the operand is in the overflow format 74;

[064] (iii) a comparator 22A, 22B that generates an asserted signal if the bits  $e_{msb} \dots e_{lsb+1}$  of the exponent field of the operand in respective buffer 11A, 11B are all binary one's and the bit  $e_{lsb}$  is either a binary one or a binary zero, which will be the case if the operand is in the overflow format 74, infinity format 75 or the NaN format 76;

[065] (iv) a comparator 23A, 23B that generates an asserted signal if the bit  $e_{msb}$  of the exponent field of the operand in respective buffer 11A, 11B is a binary one and respective bits  $e_{msb-1} \dots e_{lsb}$  are either binary one or a binary zero, which may be the case if the operand is in the normalized non-zero format 73 and will be the case for the overflow format 74, infinity format 75 or NaN format 76;

[066] (v) a comparator 24A, 24B that generates an asserted signal if the bit  $e_{msb}$  of the exponent field of the operand in respective buffer 11A, 11B is a binary zero and bits  $e_{msb-1} \dots e_{lsb}$  are all binary one's, which may be the case if the operand is in the normalized non-zero format 73 and will be the case if the operand has the value +1.0 or -1.0;

[067] (vi) a comparator 25A, 25B that generates an asserted signal if the bit  $e_{msb}$  of the exponent field of the operand in respective buffer 11A, 11B is a binary zero and respective bits  $e_{msb-1} \dots e_{lsb}$  are either binary one or binary zero, which will be

[072] (xi) a comparator 33A, 33B that generates an asserted signal if the bits  $f_{lsb+4} \dots f_{lsb+l}$  of the fraction field of the operand in the respective buffer 11A, 11B are

[073] (xii) a comparator 34A, 34B that generates an asserted signal if the bits  $f_{lsb+4} \dots f_{lsb+1}$  of the fraction field of the operand in the respective buffer 11A, 11B are binary zero's and if the bit  $f_{lsb}$  of the fraction field is a binary "one," which will be the case if the operand is in the underflow format 71 and which may be the case if the operand is in the denormalized format 72, normalized non-zero format 73, overflow format 74, or if the flags "n," "o," "u," and "z" are clear and the flag "x" is set in the infinity format 75 or NaN format 76; and

[074] (xiii) a comparator 35A, 35B that generates an asserted signal if all of the bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the operand in the respective buffer 11A, 11B are binary zero's, which will be the case if the operand is in the zero format 70, and which may be the case if the operand is in the denormalized format 72, normalized non-zero format 73, overflow format 74, or if the flags "n," "o," "u," "z" and "x" are clear in the infinity format 75 or NaN format 76.

[075] In the illustrated embodiment, each exemplary operand analysis circuit 12A, 12B also includes combinatorial logic elements that receive selected ones of the signals from the comparators and generate characteristic signals to provide

[076] (xiv) an AND gate 50A, 50B, which will generate an asserted signal if comparators 31A, 31B, and 35A, 35B are both generating asserted signals, which will be the case if the bits  $f_{msb} \dots f_{lsb}$  of the fraction field of the operand in the respective operand buffer 11A, 11B have the bit pattern

000000000000000000000000;

[078] (xvi) an AND gate 52A, 52B, which will generate an asserted signal if comparators 30A, 30B, and 32A, 32B are both generating asserted signals, which will be the case if the bits  $f_{msb} \dots f_{lsb}$  of the fraction field of the operand in the respective operand buffer 11A, 11B have the bit pattern

111111111111111111111111;

24



otherwise stated, AND gate 40A, 40B will generate an asserted signal if the value of the bits  $f_{msb} \dots f_{lsb}$  of the fraction field of the operand in the respective operand buffer 11A, 11B has the bit pattern 000000000000000000000001 or the bit pattern 000000000000000000000000;

[080] (xviii) a NAND gate 41A, 41B that generates an asserted signal if the signal generated by comparator 26A, 26B is asserted and the signal generated by AND gate 40A and 40B is negated, which will be the case if the respective operand is in the denormalized format 72; otherwise stated, NAND gate 41A, 41B will generate an asserted signal if the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B have the pattern 00000000 and a bit of the fraction field, other than the low order bit  $f_{lsb}$ , is a "one";

[081] (xix) a NAND gate 42A, 42B that generates an asserted signal if the signal generated by comparator 24A, 24B is negated and the signal generated by comparator 25A, 25B is asserted and the signal generated by comparator 26A, 26B is negated, which will be the case if the respective operand is in the normalized format 73 and its magnitude is less than 1.0; otherwise stated, NAND gate 42A, 42B will generate an asserted signal if the high-order bit  $e_{msb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B is a "zero" but the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B do not have the bit pattern 00000000 or 01111111;

[082] (xx) an OR gate 43A, 43B that generates an asserted signal if either NAND gate 41A, 41B or NAND gate 42A, 42B is asserted, which will be the case if the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in the respective operand

buffer 11A, 11B have the pattern 00000000 and a bit of the fraction field, other than the low order bit  $f_{lsb}$ , is a "one" (NAND gate 41A, 41B) or if the high-order bit  $e_{msb}$  of the exponent field is a "zero", and the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B do not have the bit pattern 00000000 or 01111111 (NAND gate 42A, 42B);

[083] (xxi) an AND gate 44A, 44B that generates an asserted signal if comparator 24A, 24B is generating an asserted signal and AND gate 50A, 50B is generating an asserted signal, which will be the case if the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in the same operand buffer 11A, 11B have the bit pattern 01111111 and the bits  $f_{msb} \dots f_{lsb}$  of the fraction field of the operand in the same operand buffer 11A, 11B have the bit pattern 000000000000000000000000;

[084] (xxii) a NAND gate 45A, 45B that generates an asserted signal if comparator 24A, 24B is generating an asserted signal and AND gate 50A, 50B is generating a negated signal, which will be the case if the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B have the bit pattern 01111111 and at least one bit  $f_{msb} \dots f_{lsb}$  of the fraction field of the operand in the respective operand buffer 11A, 11B is "one";

[085] (xxiii) a NAND gate 46A, 46B that generates an asserted signal if comparator 23A, 23B is generating an asserted signal and comparator 22A, 22B is generating a negated signal, which will be the case if the bits  $e_{msb} \dots e_{lsb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B does not have either bit pattern 11111110 or 11111111, but the high-order bit  $e_{msb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B is "one";

[086] (xxiv) a NAND gate 47A, 47B, which will generate an asserted signal if comparator 21A, 21B is generating an asserted signal and AND gate 52A, 52B is generating a negated signal which will be the case if the bits  $e_{msb}...e_{lsb}$  of the exponent field of the operand in the respective operand buffer 11A, 11B have the bit pattern 11111110 and not every bit  $f_{msb}...f_{lsb}$  of the fraction field of the operand in the respective operand buffer 11A, 11B is "one"; and

[087] (xxv) an OR gate 48A, 48B, which will generate an asserted signal if one or more of NAND gate 45A, 45B and NAND gate 46A, 46B and NAND gate 47A, 47B generates an asserted signal.

[088] In addition, the combinatorial logic in the illustrated embodiment includes a comparator 53 that generates an asserted signal if the bits  $f_{msb}...f_{lsb+5}$  of the fraction field of the operand in operand buffer 11A represent a binary-encoded value that is larger than the binary-encoded value represented by bits  $f_{msb}...f_{lsb+5}$  of the fraction field of the operand in operand buffer 11B. The combinatorial logic also includes an XOR gate 54 that generates an asserted signal if the sign bits "s" of the operands in the operand buffers 11A and 11B have different values.

[089] Each exemplary operand analysis circuit 12A, 12B provides signals to the multiplier decision table logic 15 as shown in the following table:

- (a) the signal generated by comparator 26A, 26B
- (b) the signal generated by the comparator 21A, 21B
- (c) the signal generated by the comparator 20A, 20B
- (d) the signal generated by comparator 31A, 31B
- (e) the signal generated by AND gate 50A, 50B
- (f) the signal generated by AND gate 51A, 51B
- (g) the signal generated by AND gate 52A, 52B

- (h) the signal generated by the OR gate 43A, 43B
- (i) the signal generated by the AND gate 44A, 44B
- (j) the signal generated by the OR gate 48A, 48B

[090] In addition, the signal generated by comparator 53 is provided to the multiplier decision table logic 15, as are signals from rounding mode store 16 representative of the rounding mode and the signal generated by the XOR gate 54.

[091] In addition, the multiplier core 13 generates an overflow signal and an underflow signal, which are also provided to the multiplier decision table logic 15. The underflow signal is asserted if and only if the mathematical product should be non zero but the multiplier core produces a zero result (those skilled in the art will appreciate that this differs from the IEEE 754 definition of underflow). In the illustrated embodiment, the multiplier decision table logic 15 generates control signals for controlling the result assembler 14. In addition, control signals generated by the multiplier decision table logic 15 control sets of XOR gates 67A, 67B; 68A, 68B that control toggling of one or both of the two least significant bits  $f_{lsb}$ ,  $f_{lsb+1}$  of the fraction field of the operands before they are provided to the multiplier core 13. As noted above, the result assembler 14 receives information from a number of sources, including the XOR gate 54, the operand buffers 11A and 11B, multiplier core 13 and several predetermined value stores as described below.

[092] Under control of control signals from the multiplier decision table logic circuit 15, the result assembler 14 assembles the appropriate results representing a product of the two operands, onto a result bus 17. In general, the result assembler 14 essentially assembles the result in four segments, including a sign segment that represents the sign bit of the result, an exponent segment that represents the

exponent field of the result, a high-order fraction segment that represents the bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the result, and a low-order fraction segment that represents the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the result. It will be appreciated that the low-order fraction segment, in results in the infinity format 75 and NaN format 76, corresponds to the flags "n," "o," "u," "z", and "x". One or more of these segments will represent an embedded resulting status of the resulting floating point operand.

[093] In the illustrated embodiment, the result assembler 14 may include four elements, including a link from XOR gate 54, an exponent field selector 60, a high-order fraction field selector 61 and low-order fraction field combiner 62. The link from XOR gate 54 provides the sign of the result. As is conventional, the sign of the result is positive, in which case the sign bit would have the value "zero," if the sign bits of the operands in both operand buffers 11A, 11B are the same. On the other hand, the sign of the result is negative, in which case the sign bit would have the value "one," if the sign bits of the operands in the operand buffers 11A, 11B differ. It will be appreciated that the XOR gate 54 generates a signal of the appropriate value for the sign bit for the result, which signal is coupled onto the result bus 17.

[094] As noted above, the exemplary selector 60 couples exponent value signals representative of the exponent field of the result to the result bus 17. In this embodiment, the selector 60 receives four sets of exponent field value signals, namely, the signals from the multiplier core 13 associated with the exponent field as well as three sets of signals representative of three predetermined exponent field bit

patterns, such as the patterns depicted in FIG. 1. It will be appreciated that these predetermined exponent field bit patterns correspond to the exponent fields associated with the zero format 70, underflow format 71, overflow format 74, infinity format 75 and NaN format 76. In addition, the selector 60 receives four exponent field control signals from the multiplier decision table logic 15. One of these control signals is associated with the signals from the multiplier core 13 that are associated with the exponent field. The rest of these control signals are associated with each of the sets of exponent field value signals, respectively. In enabling the result assembler 14 to assemble the result, the multiplier decision table logic 15 will assert one of the four exponent field control signals. Additionally, the selector 60 will couple the set of exponent field value signals associated with the asserted exponent field control signal to the result bus 17 to provide the exponent field of the result.

[095] The exemplary selector 61 couples high-order fraction field signals representative of the high-order fraction field bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the result to the result bus 17. In the illustrated embodiment, the selector 61 receives seven sets of high-order fraction field value signals, namely, the signals from the multiplier core 13 associated with the high-order fraction field, signals representative of bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the operand in buffer 11A, signals representative of bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the operand in buffer 11B, as well as four sets of signals representative of four predetermined high-order fraction field bit patterns, such as the patterns depicted in FIG. 1. It will be appreciated that these predetermined high-order fraction field bit patterns correspond to the high-order fraction fields associated with the zero format 70, underflow format 71,

overflow format 74, infinity format 75 and NaN format 76. In addition, the selector 61 receives seven high-order fraction field control signals from the multiplier decision table logic 15. One of these control signals is associated with the multiplier core 13. The rest of these control signals are associated with each of the sets of high-order fraction field value signals, respectively. It will be appreciated that the control signal associated with the multiplier core 13 is the same control signal that controls the selector 60. In enabling the result assembler 14 to assemble the result, the multiplier decision table logic will assert one of the seven high-order fraction field control signals, and the selector 61 will couple the set of high-order fraction field value signals associated with the asserted high-order fraction field control signal to the result bus 17 to provide bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the result.

[096] Similarly, the exemplary combiner 62 couples low-order fraction field value signals representative of the low-order fraction field bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result to the result bus 17. The combiner 62 receives four sets of low-order fraction field signals, namely, the signals from the multiplier core 13 associated with the low-order fraction field, signals representative of bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the operand in buffer 11A, signals representative of bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the operand in buffer 11B, and one set of signals from the multiplier decision table logic 15. It will be appreciated that in the illustrated embodiment, the set of signals provided by the multiplier decision table logic 15 are used in controlling the condition of flags "n", "o", "u", "z", and "x" for those formats in which the low order bits  $f_{lsb+4} \dots f_{lsb}$  represent flags. In addition, the sets of signals provided by the operands in buffers 11A and 11B may also represent the flags "n",

"o", "u", "z", and "x". In addition, the combiner 62 receives three low-order fraction field control signals from the multiplier decision table logic 15. One control signal is associated with the sets of low-order fraction field value signals provided by the multiplier core and the two others are associated with the sets of signals provided by the buffers 11A and 11B.

[097] In enabling the result assembler 14 to assemble the result, the multiplier decision table logic 15 may provide signals representative of the low-order fraction field and negate all of the low-order fraction field control signals. When this occurs, the signals representative of the low order fraction field provided by the multiplier decision table logic 15 will be coupled to the result bus 13 to provide bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result.

[098] Alternatively, the multiplier decision table logic 15 may negate all of the low-order fraction field value signals provided thereby and assert one of the three low-order fraction field control signals. When this occurs, the combiner 62 will couple the set of low-order fraction field value signals associated with the asserted low-order fraction field control signal to the result bus 17 to provide bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result. As a further alternative, the multiplier decision table logic 15 may negate all of the low-order fraction field value signals provided thereby and assert more than one of the three low-order fraction field control signals. As a result, the combiner 62 will couple the bit-wise OR of the sets of low-order fraction field value signals associated with the asserted low-order fraction field control signals to the result bus 17 to provide bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result. As yet another alternative, the multiplier decision table logic 15 may assert one or



[099] In more detail, the exemplary combiner 62 in the illustrated embodiment comprises an OR circuit 63 and three AND circuits 64 through 66. (Each gate in the diagram actually represents five such gates, one for each bit position  $f_{lsb+4} \dots f_{lsb}$ , but for the sake of clarity and to avoid confusion, the diagram illustrates them as a single gate.) The AND circuits 64-66 receive the low-order fraction field value signals from the multiplier core 13 and operand buffers 11A and 11B respectively, as well as the respective low-order fraction field control signal. These AND circuits 64-66 perform a bit-wise AND operation to, if the respective low-order fraction field control signal is asserted, couple the low-order fraction field value signals to a respective input of OR circuit 63. The OR circuit 63, whose output is connected to the result bus 17, performs a bit-wise OR operation in connection with the signals that it receives from the AND circuits 64-66 and the low-order fraction field value signals provided by the multiplier decision table logic 15. If the multiplier decision table logic 15 negates all of the low-order fraction field control signals, the AND circuits 64-66 will block the low-order fraction field value signals that they receive, and the signals provided by the OR circuit 63 will conform to the low-order fraction field value signals provided by the multiplier decision table logic 15.

[0100] On the other hand, if the multiplier decision table logic 15 asserts one or more of the low-order fraction field control signals, the AND circuits 64-66 that receive the asserted low-order fraction field control signal will couple the low-order fraction field value signals that they receive to the OR circuit 63 and the other AND gates will block the low-order fraction field signal that they receive. As will be described below, under some circumstances, the multiplier decision table logic 15 will assert two low-order fraction field control signals to enable two sets of low-order fraction field value signals to be coupled to the OR circuit 63. In that case, the OR gate will perform a bit-wise OR operation in connection with signals representing respective bits of the low-order fraction field.

[0101] Thus, multiplier decision table logic 15 will assert two low-order fraction signals if, for example, both operands in operand buffers 11A and 11B are in NaN format to enable the respective flags "n," "o," "u," and "x" to be ORed together. However, if the low-order fraction field value signals provided by the multiplier decision table logic 15 are negated, the low-order fraction field value signals provided by the OR circuit 63 will conform to the low-order fraction field signals provided by the AND circuit or circuits that receive the asserted low-order fraction field control signal.

[0102] As noted above, the multiplier decision table logic 15 generates control signals for controlling the selectors 60 and 61 and combiner 62 comprising the result assembler 14 and for controlling the toggling of the signals representing the low-order bits  $f_{lsb+1}$  and  $f_{lsb}$  from the operand buffers 11A and 11B before they are presented to the multiplier core 13. The control signals generated by the multiplier

decision table logic 15 are such as to enable the result to be assembled in the desired format 70-76 having status information embedded within the result itself. Before proceeding further, it would be helpful to describe the results that are to be generated by the multiplier unit 10.

[0103] Generally, exemplary results generated by the multiplier unit 10 are described in the table depicted in FIG. 3. In that table, one skilled in the art will appreciate that "+P" or "+Q" means any finite positive representable value greater than "one," other than +OV (that is, a value in the overflow format 74 with the sign bit "s" being "zero"). "-P" or "-Q" means any finite negative representable value less than negative-one, other than -OV (that is, a value in the overflow format 74, with the sign bit being "one"). "+R" or "+S" means any positive non-zero representable value less than "one," other than +UN (that is, a value in the underflow format 71 with the sign bit "s" being "zero"). "-R" or "-S" means any negative non-zero representable value greater than negative-one, other than -UN (that is, a value in the underflow format 71, with the sign bit being "one"). Finally, those skilled in the art will appreciate that "NaN" means any value whose exponent field is 11111111, other than one of the values represented by  $+\infty$  (that is, a value in the infinity format 75, with the sign bit "s" being "zero") and  $-\infty$  (that is, a value in the infinity format 75, with the sign bit "s" being "one").

[0104] Key to symbols in the table with exemplary results depicted in FIG. 3 are as follows:

[0105] (a) The result is  $+\infty$ , with the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of

[0106] (b) The result is  $+\infty$ , with the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of

[0107] (c) The result is  $+\infty$ , with the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of

[0108] (d) The result is  $+\infty$ , with five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the

[0109] (e) For "round toward plus infinity", the result is  $+\infty$ , with the five

[0110] (f) For "round toward plus infinity," the result is  $-0$ . For "round toward minus infinity," the result is  $-\infty$ , with five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field being equal to the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the infinite operand. For all other rounding modes, the result is a negative NaN value 1 11111111 1000000000000001001ouzx (to indicate "zero times infinity" with the invalid operation flag set), where ouzx are the four least significant bits  $f_{lsb+3} \dots f_{lsb}$  of the fraction field of the infinite operand.

[0111] (g) The result is  $-\infty$ , with the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result being the bitwise OR of the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the infinite operand with 00101 (to indicate underflow and inexact).

[0112] (h) The result is  $-\infty$ , with the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field being equal to the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the infinite operand.

[0113] (i) The result is  $-\infty$ , with the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result being the bitwise OR of the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the infinite operand with 01001 (to indicate overflow and inexact).

[0114] (j) The result is  $-\infty$ , with five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result being the bitwise OR of the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction fields of the operands.

[0115] (k) The result is a copy of the NaN operand, except that its sign is reversed if the other operand is negative, and that the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result are the bitwise OR of the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction fields of the operands.

[0116] (l) For "round toward plus infinity," the result is the same as if -OV were replaced by -HUGE. For all other rounding modes, the result is +OV.

[0117] (m) For "round toward plus infinity," the result is +OV. For "round toward minus infinity," the result is +UN. For all other rounding modes, the result is the positive NaN value 0 11111111 10000000000000010111101 (to indicate "UN times OV" with the invalid operation, overflow, underflow, and inexact flags set).

[0118] (n) For "round toward plus infinity," the result is -UN. For "round toward minus infinity," the result is -OV. For all other rounding modes, the result is the negative NaN value 1 11111111 10000000000000010111101 (to indicate "UN times OV" with the invalid operation, overflow, underflow, and inexact flags set).

[0119] (o) For "round toward minus infinity," the result is the same as if -OV were replaced by -HUGE. For all other rounding modes, the result is -OV.

[0120] (p) The result is a copy of the NaN operand, except that its sign is reversed if the other operand is negative. The five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result are ORed with 01001 (to indicate overflow and inexact).

[0121] (q) The result is as computed in accordance with IEEE Std. 754. However, if overflow occurs or if the rounding mode is "round toward plus infinity" and the mathematical product is greater than +HUGE, the result is +OV. Further, if

[0122] (r) For "round toward plus infinity," the result is the same as if -UN were replaced by -TINY. For all other rounding modes, the result is as computed in accordance with IEEE Std. 754.

[0124] (t) The result is as computed in accordance with IEEE Std. 754. However, if overflow occurs or if the rounding mode is "round toward minus infinity" and the mathematical product is less than -HUGE, the result is -OV. Further, if underflow occurs and a computation in accordance with IEEE Std 754 would provide the result -0 or if the rounding mode is "round toward plus infinity" and the mathematical product is greater than -TINY, the result is -UN.

[0126] (v) For "round toward minus infinity," the result is the same as if +OV were replaced by +HUGE. For all other rounding modes, the result is -OV.

39

[0128] (x) The result is a copy of the NaN operand, except that its sign is reversed if the other operand is negative. The five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result are ORed with 00101 (to indicate underflow and inexact).

[0129] (y) For "round toward plus infinity," the result is the same as if +UN were replaced by +TINY. For all other rounding modes, the result is as computed in accordance with IEEE Std. 754.

[0130] (z) For "round toward minus infinity," the result is the same as if +OV were replaced by +HUGE. For all other rounding modes, the result is +OV.

[0131] (@) The result is a copy of the NaN operand that has the larger value in the fraction field, except that the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result are the bitwise OR of the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the operands and the sign bit of the result is 1, indicating a negative result, if and only if the sign bits of the two NaN operands differ.

[0132] Multiplication is commutative with multiplier unit 10 operating according to the table depicted in FIG. 3. This is true even with those cases where one or both operands are values in the NaN format 76.

[0133] As noted above, multiplier decision table logic 15 generates control signals for controlling the selectors 60, 61 and combiner 62 comprising the result assembler 14 and for controlling the toggling of the signals representing the low-order bits  $f_{lsb+1}$  and  $f_{lsb}$  from the operand buffers 11A and 11B before they are presented to the multiplier core 13. The particular signals that the multiplier decision table logic 15 will generate depends on the signals provided thereto by the operand



buffers 11A and 11B representing the states of the respective sign bits, the operand analysis circuits 12A and 12B, comparator 40, rounding mode store 16, and the overflow and underflow signals from the multiplier core 13. In the illustrated embodiment, the series of input signals received by the multiplier decision table logic 15 are as follows:

[0134] (a) a signal from comparator 26A that is asserted if the exponent field of the operand in operand buffer 11A has the bit pattern 00000000;

[0135] (b) a signal from comparator 21A that is asserted if the exponent field of the operand in operand buffer 11A has the bit pattern 11111110;

[0136] (c) a signal from comparator 20A that is asserted if the exponent field of the operand in operand buffer 11A has the bit pattern 11111111;

[0137] (d) a signal from the comparator 31A that is asserted if the operand in operand buffer 11A has a high order fraction field with all 0-bits;

[0138] (e) a signal from AND gate 50A that is asserted if the operand in operand buffer 11A has high- and low-order fraction fields with the collective bit pattern 000000000000000000000000;

[0139] (f) a signal from AND gate 51A that is asserted if the operand in operand buffer 11A has high- and low-order fraction fields with the collective bit pattern 000000000000000000000001;

[0140] (g) a signal from AND gate 52A that is asserted if the operand in operand buffer 11A has high- and low-order fraction fields with the collective bit pattern 111111111111111111111111;

[0141] (h) a signal from OR gate 43A that is asserted if any of the following signals are asserted:

[0142] (1) a signal from NAND gate 41A that is asserted if the exponent field of the operand in operand buffer 11A has a bit pattern 00000000 (which will be the case if the signal from comparator 26A is asserted) and the high- and low-order fraction field of the operand in operand buffer 11A has a bit pattern in which at least one bit other than the least significant bit, is "1" (which will be the case if the signal from AND gate 40A is negated); and

[0143] (2) a signal from NAND gate 42A that is asserted if the exponent field of the operand in operand buffer 11A does not have a bit pattern 01111111 (which will be the case if the signal from comparator 24A is negated) and also does not have the bit pattern 00000000 (which will be the case if the signal from comparator 26A is negated), but the high-order bit  $e_{msb}$  of the exponent field in operand buffer 11A is "0" (which will be the case if the signal from comparator 25A is asserted);

[0144] (i) a signal from AND gate 44A that is asserted if the exponent field of the operand in operand buffer 11A has the bit pattern 01111111 (which will be the case if the signal from comparator 24A is asserted) and the bits of the fraction field of the operand in operand buffer 11A are all "0" (which, will be the case if the signal from AND gate 50A is asserted);

[0145] (j) a signal from OR gate 48A that is asserted if any of the following signals are asserted:

[0146] (1) a signal from NAND gate 45A that is asserted if the exponent field of the operand in operand buffer 11A has the bit pattern 01111111 (which will be the case if the signal from comparator 24A is asserted) and at least one bit of the fraction field of the operand in operand buffer 11A is "1" (which will be the case if the signal from AND gate 50A is negated);

[0147] (2) a signal from NAND gate 46A that asserted if the exponent field of the operand in operand buffer 11A does not have the bit pattern 11111110 or 11111111 but the high-order bit of the exponent field is 1 (which will be the case if the signal from comparator 23A is asserted and the signal from comparator 22A is negated); and

[0148] (3) a signal from NAND gate 47A that is asserted if the exponent field of the operand in operand buffer 11A has the bit pattern 11111110 (which will be the case if the signal from comparator 21A is asserted) and at least one bit of the fraction in the fraction field of the operand in operand buffer 11A is "0" (which will be the case if the signal from AND gate 52A is negated);

[0149] (k) a signal from comparator 26B that is asserted if the exponent field of the operand in operand buffer 11B has the bit pattern 00000000;

[0150] (l) a signal from comparator 21B that is asserted if the exponent field of the operand in operand buffer 11B has the bit pattern 11111110;

[0151] (m) a signal from comparator 20B that is asserted if the exponent field of the operand in operand buffer 11B has the bit pattern 11111111;

[0152] (n) a signal from the comparator 31B that is asserted if the operand in operand buffer 11B has a high-order fraction field with all 0-bits;

[0153] (o) a signal from AND gate 50B that is asserted if the operand in operand buffer 11B has high- and low-order fraction fields with the collective bit pattern 000000000000000000000000;

[0154] (p) a signal from AND gate 51 B that is asserted if the operand in operand buffer 11B has high- and low-order fraction fields with the collective bit pattern 000000000000000000000001;

[0155] (q) a signal from AND gate 52B that is asserted if the operand in operand buffer 11B has high- and low-order fraction fields with the collective bit pattern 111111111111111111111111;

[0156] (r) a signal from OR gate 43B that is asserted if any of the following signals are asserted:

[0157] (1) a signal from NAND gate 41B that is asserted if the exponent field of the operand in operand buffer 11B has a bit pattern 00000000 (which will be the case if the signal from comparator 26B is asserted) and the high- and low-order fraction field of the operand in operand buffer 11B has a bit pattern in which at least one bit, other than the least significant bit, is "1" (which will be the case if the signal from AND gate 40B is negated); and

[0158] (2) a signal from NAND gate 42B that is asserted if the exponent field of the operand in operand buffer 11B does not have a bit pattern 01111111 (which will be the case if the signal from comparator 24B is negated) and also does not have the bit pattern 00000000 (which will be the case if the signal from comparator 26B is negated), but the high-order bit  $e_{msb}$  of the exponent field in

[0159] (s) a signal from AND gate 44B that is asserted if the exponent field of the operand in operand buffer 11B has the bit pattern 01111111 (which will be the case if the signal from comparator 24B is asserted) and the bits of the fraction field of the operand in operand buffer 11B are all "0" (which will be the case if the signal from AND gate 50B is asserted);

[0161] (1) a signal from NAND gate 45B that is asserted if the exponent field of the operand in operand buffer 11B has the bit pattern 01111111 (which will be the case if the signal from comparator 24B is asserted) and at least one bit of the fraction field of the operand in operand buffer 11B is "1" (which will be the case if the signal from AND gate 50B is negated);

[0163] (3) a signal from NAND gate 47B that is asserted if the exponent field of the operand in operand buffer 11B has the bit pattern 11111110 (which will be the case if the signal from comparator 21B is asserted) and at least



[0173] (3) a signal that, if asserted, enables the second-least significant bit of the operand in operand buffer 11B to be toggled before being presented to the multiplier core 13;

[0174] (4) a signal that, if asserted, enables the least significant bit of the operand in operand buffer 11B to be toggled before being presented to the multiplier core 13;

[0175] (5) a signal that, if asserted, enables the exponent field and the high-part of the fraction of the result to be provided by the multiplier core 13, and moreover, the five least-significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the output provided by the multiplier core 13 will contribute to the five least significant bits  $f_{lsb+4} \dots f_{lsb}$  of the result;

[0176] (6) a signal that, if asserted, will enable the exponent field of the result to have the bit pattern 00000000;

[0177] (7) a signal that, if asserted, will enable the exponent field of the result to have the bit pattern 11111110;

[0178] (8) a signal that, if asserted, will enable the exponent field of the result to have the bit pattern 11111111;

[0179] (9) a signal that, if asserted, will enable the high-order fraction of the result to correspond to the high-order portion of the fraction of the operand in operand buffer 11A;

[0180] (10) a signal that, if asserted, will enable the high-order fraction of the result to correspond to the high-order portion of the fraction of the operand in operand buffer 11B;

[0181] (11) a signal that, if asserted, will enable the high-order fraction of the result to correspond to the bit pattern 000000000000000000;

[0182] (12) a signal that, if asserted, will enable the high-order fraction of the result to correspond to the bit pattern 111111111111111111;

[0183] (13) a signal that, if asserted, will enable the high-order fraction of the result to correspond to the bit pattern 100000000000000100 (to represent a NaN value "zero times infinity");

[0184] (14) a signal that, if asserted, will enable the high-order fraction of the result to correspond to the bit pattern 100000000000000101 (to represent a NaN value "underflow times overflow");

[0185] (15) a signal that, if asserted, will enable the low-order fraction field of the operand in output buffer 11A to contribute to the five least-significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result;

[0186] (16) a signal that, if asserted, will enable the low-order fraction field of the operand in output buffer 11B to contribute to the five least-significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result; and

[0187] (17)-(21) signals that always contribute to the five least-significant bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction field of the result.

[0188] The specific patterns of output signals (1) through (21) generated by the exemplary multiplier decision table logic 15 in response to patterns of input signals (a) through (z) are depicted in FIGS. 4A through 4C. Generally, in FIGS. 4A-4C, each row represents conditions of the output signals (1) through (21) that are generated by the multiplier decision table logic 15 in response to one pattern of



input signals (a) through (z). In each row, the indicia to the left of the asterisk (\*) represent the pattern of input signals (a) through (z) and the indicia to the right of the asterisk represent the pattern of output signals (1) through (21) with a "1" indicating that the respective input or output signal is asserted, a "0" indicating that the respective input or output signal is negated and a "-" indicating that the respective input signal may be either negated or asserted. Each row is further annotated with an indication as to the respective format 70 through 76 of the operand in the respective operand buffers 11A and 11B and the format of the result.

[0189] Referring now to FIG. 4A, a discussion of the first row of input signal values and corresponding output signal values follows.

[0190] (A) for the three input patterns to the left of the asterisk:

[0191] (i) the first pattern "--10-----" indicates that signal (c) is asserted, signal (d) is negated, and signals (a), (b), and (e) through (j) may be either asserted or negated, with the pattern indicating a value in the NaN format 76 ("[NaN]");

[0192] (ii) the second pattern "--10-----" indicates that signal (m) is asserted, signal (n) is negated, and signals (k), (l) and (o) through (t) may be either asserted or negated, with the pattern indicating a value in the NaN format 76 ("[NaN]"); and

[0193] (iii) the third pattern "1 -- - --" indicates that the signal (u) is asserted, indicating that the binary-encoded value of the high-order bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the operand in operand buffer 11A is greater than the binary-encoded value of the high-order bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the

operand in operand buffer 11B, and signals (v) through (z) may be either asserted or negated; and

[0194] (B) for the six output patterns to the right of the asterisk:

[0195] (i) the pattern "0000" to the immediate right of the asterisk indicates that the signals provided to XOR gates 67A, 68A, 67B and 68B, which control the toggling of the low-order bits  $f_{lsb+1}$  and  $f_{lsb}$  of the fraction fields of the operands in operand buffers 11A and 11B before being presented to multiplier core 13, are all negated;

[0196] (ii) the next "0" indicates that the signal provided to selectors 60 and 61 and AND circuit 64, is negated thereby to ensure that the output from multiplier core 13 will not contribute the result;

[0197] (iii) the next pattern "001" indicates that the signal will be asserted that will enable selector 60 to couple signals representative of the pattern 11111111 to the result bus 17, and the signals associated with the other patterns 00000000 and 11111110 will be negated;

[0198] (iv) the next pattern "100000" indicates selector 61 is enabled to couple the signals associated with the bits  $f_{msb}...f_{lsb+5}$  (comprising the high-order fraction field of the operand in buffer 11A) to the result bus 17 as the bits  $f_{msb}...f_{lsb+5}$  of the fraction field of the result;

[0199] (v) the next pattern "11" indicates that the signals provided by the multiplier decision table logic 15 will provide asserted signals to both AND circuits 65 and 66, enabling both AND gates to couple signals received thereby from both operand buffers 11A and 11B to the OR circuit 63; and

[0200] (vi) the last pattern "00000" indicates that the signals provided by the multiplier decision table logic 15 to the OR circuit 63 are all negated. With such a last pattern, the OR gate will perform a bit-wise OR operation in connection with those signals and the signals provided thereto by AND circuits 65 and 66. Further, the negated signal described in (B)(ii) provides that the signals provided by AND circuit 64 are also negated in which case the signals coupled by OR circuit 63 to result bus 17 will correspond to the OR of the bits  $f_{lsb+4} \dots f_{lsb}$  from the fraction fields of the operands in operand buffers 11A and 11B.

[0201] On the right hand side of the first row in FIG. 4A, the legend "[NaN op1 f1|f2]" indicates that the result value is in the NaN format 76 with the bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the result corresponding to bits  $f_{msb} \dots f_{lsb+5}$  of the fraction field of the operand in operand buffer 11A and the bits  $f_{lsb+4} \dots f_{lsb}$  of the result corresponding to the OR of the bits  $f_{lsb+4} \dots f_{lsb}$  of the fraction fields of the operands in both operand buffers 11A and 11B. It should be noted that this corresponds to the result represented by symbol (@) in the table depicted in FIG. 3.

[0202] In the context of the above discussion, the other rows of FIG. 4A and the rows in FIGS. 4B-4C will be apparent to those skilled in the art.

[0203] As described above, the sign of the result is provided directly by the XOR gate 54.

[0204] Multiplier decision table logic 15 may be implemented by many different circuit elements that will be apparent to those skilled in the art, including, but not limited to programmable logic arrays, ASIC circuits, general memory registers, other addressable memory storage devices or a combination thereof.

[0205] One of ordinary skill in the art will recognize that other formats and bit patterns could be used to represent the floating point operand formats without departing from the principles of the present invention. One of ordinary skill in the art will also recognize that the floating point status information contained in the operands could easily be represented by other bit combinations (not shown) without departing from the principles of the present invention. For example, more or fewer bits could be used, a subset or superset of the exemplary status bits could be used, or the most significant bits of an operand (or some other subset of bits) could be used to indicate the floating point status information, instead of the least significant bits illustrated.

[0206] It will be appreciated that a system in accordance with an embodiment of the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof. Any portion of such a system may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided in to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

[0207] The foregoing description has been limited to a specific embodiment of this invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention.

[0208] Other embodiments of the invention will be apparent to those skilled in the art from consideration of the specification and practice of the invention disclosed herein. It is intended that the specification and examples be considered as exemplary only, with a true scope and spirit of the invention being indicated by the following claims.

FOOTNOTES